Simple set cardinality estimation through random sampling

Marco Bressan
Dip. Informatica
Sapienza Università di Roma
Roma, Italy
bressan@di.uniroma1.it

Enoch Peserico
Dip. Ing. Informazione
Università di Padova
Padova, Italy
enoch@dei.unipd.it

Luca Pretto
Dip. Ing. Informazione
Università di Padova
Padova, Italy
pretto@dei.unipd.it

arXiv:1512.07901v2 [cs.DM] 18 Jan 2016

Abstract

We present a simple algorithm for estimating the cardinality of a set I, based on a $RandomSample(I)$ primitive that returns an element of I uniformly at random. Our algorithm with probability $(1-p_{err})$ returns an estimate of $|I|$ accurate within a factor $(1 \pm \delta_{err})$ invoking $RandomSample(I)$ at most $O\big(\frac{1}{\delta_{err}}\sqrt{\log(\frac{1}{p_{err}})}\sqrt{|I|}\big)$ times (for $\frac{1}{\delta_{err}}\sqrt{\log(\frac{1}{p_{err}})} = O(\sqrt{|I|})$).

Tasks like graph size estimation (see e.g. [1] and [2]) have recently revived interest in the problem of estimating the cardinality of a set via random sampling. This short note gives a simple algorithm for estimating the size of a set within a desired precision with a given probability of error. Although the estimator we use has long been known (see [3] and [4]), our result adds guarantees on the number of samples needed parameterized on precision and error probability, and can be readily used "black box" by other algorithms.

1 Estimating set cardinality via random sampling

We present below the pseudocode of an algorithm for estimating the cardinality of a set I based on repeated invocations of a routine $RandomSample(I)$ that, on each invocation, returns an element of I chosen uniformly and independently at random. We then formally state and prove the bounds on the probability that $RandomSample(I)$ is invoked "too many" times, or that the estimate of $|I|$ is not accurate within a factor $(1 \pm \delta_{err})$. We assume hereafter without loss of generality that the required δ_{err} is less than 1 (or the trivial estimator $|I| = 0$ would suffice).

Algorithm CardinalityEstimator

$s \leftarrow 0$ ▷ # of samples taken so far
$d \leftarrow 0$ ▷ # of *distinct* samples taken so far
$w \leftarrow 0$ ▷ # of samples taken so far, each weighted by d when taken
repeat
  $w \leftarrow w + d$
  $s \leftarrow s + 1$
  $sample[s] \leftarrow RandomSample(I)$
  if $sample[s] \notin \bigcup_{i=1}^{s-1} sample[i]$ **then** $d \leftarrow d + 1$
until $(s - d) \geq \frac{4}{\delta_{err}^2} \ln\big(\frac{3}{p_{err}}\big)$
return $\frac{w}{s-d}$

Theorem. CardinalityEstimator *with probability greater than* $(1-p_{err})$ *outputs an estimate* $|\tilde{I}|$ *such that* $(1-\delta_{err})|I| \leq |\tilde{I}| \leq (1+\delta_{err})|I|$ *invoking RandomSample*(I) *at most* $\min(|I|, 2\lceil\sqrt{k_{err}|I|}\rceil) + \lceil k_{err}\rceil$ *times, where* $k_{err} = \frac{4}{\delta_{err}^2}\ln\left(\frac{3}{p_{err}}\right)$.

Proof. Suppose CardinalityEstimator has so far taken exactly $i-1$ samples, and denote by $d(i)$ and $w(i)$ the values of d and w just before the i^{th} sample is taken (i.e. just before line 7 is executed for the i^{th} time). The probability that the next invocation of *RandomSample*(I) yields a repeat is then $\frac{d(i)}{|I|}$. Therefore, although clearly $d(i)$ is determined by the outcomes of previous samples, once its value is known we can associate to the event that the next invocation of *RandomSample*(I) yields a repeat an indicator random variable $\boldsymbol{\chi}_i^{d(i)}$ with expectation $\frac{d(i)}{|I|}$; and, crucially, this random variable is independent from $\boldsymbol{\chi}_1^{d(1)}, \ldots, \boldsymbol{\chi}_{i-1}^{d(i-1)}$. Therefore, whenever line 9 is executed, the number $s-d$ of repeats witnessed so far by the algorithm is the realization of the sum of s independent indicator random variables $\boldsymbol{\chi}_1^{d(1)}, \ldots, \boldsymbol{\chi}_s^{d(s)}$ with total expectation

$$\mathbb{E}\Big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\Big] = \sum_{i=1}^{s}\frac{d(i)}{|I|} = \frac{w(s)}{|I|} \tag{1}$$

At line 10, the algorithm returns as an estimate of $|I|$ the value $|\tilde{I}|$ that would make the formula for $\mathbb{E}\big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\big]$ given by Equation 1 match the actual value $s-d$ of the realization of $\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}$, i.e. $|\tilde{I}| = \frac{w(s)}{s-d}$; for this estimate to deviate excessively from $|I|$ we must then have $\mathbb{E}\big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\big]$ far enough from the actual realization $s-d \geq k_{err}$ of $\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}$. In a nutshell, the algorithm stops acquiring samples when the sum of the realizations of the (independent) $\boldsymbol{\chi}_i^{d(i)}$ is sufficiently large to make a significant deviation from its expectation unlikely.

Let us begin by bounding the probability that the estimate exceeds $(1+\delta_{err})|I|$. For this event to take place, the algorithm must terminate with $\frac{w(s)}{s-d} = (1+\delta)|I|$ for some $\delta > \delta_{err}$. Thus $\mathbb{E}\big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\big] = \frac{w(s)}{|I|} = (1+\delta)(s-d) \geq (1+\delta)k_{err}$, and $\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}$ has taken value $s-d = (1-\frac{\delta}{1+\delta})\mathbb{E}\big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\big]$. By the bounds of Appendix A.1, the probability of this event is then at most

$$Pr\Big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)} \leq \big(1-\frac{\delta}{1+\delta}\big)\mathbb{E}\Big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\Big]\Big] \leq e^{-\frac{(1+\delta)k_{err}(\frac{\delta}{1+\delta})^2}{2}} = \Big(\frac{p_{err}}{3}\Big)^{\frac{2\delta^2}{\delta_{err}^2(1+\delta)}} \tag{2}$$

and the last term is smaller than $\frac{p_{err}}{3}$ since $\frac{2\delta^2}{\delta_{err}^2(1+\delta)} \geq \frac{\delta^2}{\delta_{err}^2} > 1$ when $\delta \leq 1$ and $\frac{2\delta^2}{\delta_{err}^2(1+\delta)} > \frac{\delta}{\delta_{err}^2} > 1$ when $\delta > 1$.

In a similar fashion we can bound the probability that the estimate falls below $(1-\delta_{err})|I|$. In this case $\frac{w(s)}{s-d} = (1-\delta)|I|$ for some $1 > \delta > \delta_{err}$, thus $\mathbb{E}\big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\big] = \frac{w(s)}{|I|} = (1-\delta)(s-d) \geq (1-\delta)k_{err}$ and $\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}$ has taken value $s-d = (1+\frac{\delta}{1-\delta})\mathbb{E}\big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\big]$. Again by the bounds of Appendix A.1, the probability of this event is no larger than

$$Pr\Big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)} \geq \big(1+\frac{\delta}{1-\delta}\big)\mathbb{E}\Big[\sum_{i=1}^{s}\boldsymbol{\chi}_i^{d(i)}\Big]\Big] \leq e^{-\frac{(1-\delta)k_{err}(\frac{\delta}{1-\delta})^2}{2+\frac{\delta}{1-\delta}}} = \Big(\frac{p_{err}}{3}\Big)^{\frac{4\delta^2}{\delta_{err}^2(2-\delta)}} \tag{3}$$

and the last term is smaller than $\frac{p_{err}}{3}$ since $\frac{4\delta^2}{\delta_{err}^2(2-\delta)} > \frac{\delta^2}{\delta_{err}^2} > 1$.

Let us now prove less than $\frac{p_{err}}{3}$ the probability that the algorithm invokes $RandomSample(I)$ more than $2\lceil\sqrt{k_{err}|I|}\rceil + k_{err}$ times before obtaining k_{err} repeats and thus terminating. Clearly, this event takes place if and only if $RandomSample(I)$ yields $2\lceil\sqrt{k_{err}|I|}\rceil + 1$ distinct samples before yielding k_{err} repeats, and thus, denoting by $\mathbf{r}(d)$ the random variable giving the total number of repeats yielded by $RandomSample(I)$ before the $(d+1)^{th}$ distinct sample is obtained, if and only if $\mathbf{r}(2\lceil\sqrt{k_{err}|I|}\rceil) < k_{err}$. Denote by $\boldsymbol{\rho}_i$ the indicator variable of the event of obtaining at least one repeat between the i^{th} and $(i+1)^{th}$ distinct samples; note that obviously $\mathbf{r}(d) \geq \sum_{i=1}^{d} \boldsymbol{\rho}_i$, that all $\boldsymbol{\rho}_i$ are independent, and that $\mathbb{E}[\boldsymbol{\rho}_i] = \frac{i}{|I|}$ and thus $\mathbb{E}\big[\sum_{i=1}^{2\lceil\sqrt{k_{err}|I|}\rceil} \boldsymbol{\rho}_i\big] = \sum_{i=1}^{2\lceil\sqrt{k_{err}|I|}\rceil} \frac{i}{|I|} > 2k_{err}$. Then, once again by the bounds of Appendix A.1,

$$Pr\Big[\mathbf{r}\big(2\lceil\sqrt{k_{err}|I|}\,\rceil\big) < k_{err}\Big] < e^{-\frac{(1/2)^2}{2}\cdot 2k_{err}} = e^{-\frac{1}{\delta_{err}^2}\ln\frac{3}{p_{err}}} < \frac{p_{err}}{3} \tag{4}$$

Noting that obviously $|I| + k_{err}$ samples always yield at least k_{err} repeats, we have that the probability that $CardinalityEstimator$ invokes $RandomSample(I)$ more than $\min(|I|, 2\lceil\sqrt{k_{err}|I|}\rceil) + k_{err}$ times is less than $\frac{p_{err}}{3}$. Since the probabilities that the estimate returned by $CardinalityEstimator$ falls below $(1-\delta_{err})|I|$ or exceeds $(1+\delta_{err})|I|$ are also each less than $\frac{p_{err}}{3}$, by a simple union bound we obtain the thesis. □

A Appendix

A.1 Probability bounds

Lemma 1 (Chernoff bounds, Angluin-Valiant version)**.** *Let X be the sum of n independent 0-1 random variables $X_1, \ldots, X_n$. Then, for any $\Delta > 0$, we have:*

$$Pr[X \leq (1-\Delta)\mathbb{E}[X]] \leq e^{-\frac{\Delta^2}{2}\mathbb{E}[X]}$$

$$Pr[X \geq (1+\Delta)\mathbb{E}[X]] \leq e^{-\frac{\Delta^2}{2+\Delta}\mathbb{E}[X]}$$

References

[1] S. J. Hardiman and L. Katzir. Estimating Clustering Coefficients and Size of Social Networks via Random Walk. *ACM Transactions on the Web*, 9(4):19:1–19:20, 2015.

[2] L. Katzir, E. Liberty, O. Somekh, I. A. Cosma. Estimating Sizes of Social Networks via Biased Sampling. *Internet Mathematics*, 10(3–4):335–359, 2014.

[3] Z. E. Schnabel. The Estimation of Total Fish Population of a Lake. *The American Mathematical Monthly*, 45(6):348–352, 1938.

[4] L. A. Goodman. Sequential Sampling Tagging for Population Size Problems. *The Annals of Mathematical Statistics*, 24(1):56–69, 1953.